MA 3046 - Matrix Analysis
Laboratory Number 2
The Singular Value Decomposition (SVD)

There are three "classic" matrix factorizations in linear algebra, the $\mathbf{L}\,\mathbf{U}$ decomposition (which you should have seen in MA1043), the $\mathbf{Q}\,\mathbf{R}$ decomposition (which we shall see later in this course), and the Singular Value decomposition (SVD). In the SVD, the matrix $\mathbf{A} \in \mathbb{C}^{m\times n}$ is decomposed (factored):

$$\mathbf{A} \;=\; \mathbf{U}\,\boldsymbol{\Sigma}\;\mathbf{V}^H$$

where $\mathbf{U}$ and $\mathbf{V}$ are unitary, and $\boldsymbol{\Sigma}$ is diagonal. The diagonal elements of $\boldsymbol{\Sigma}$ are

$$\sigma_i \;=\; \sqrt{\lambda_i}$$

where, by convention $\sigma_1 \geq \sigma_2 \geq \cdots$, and the λ_i are the (assured non-negative) eigenvalues of $\mathbf{A}^H\,\mathbf{A}$. The columns of $\mathbf{V}$ are called the *right singular vectors*, and are also the eigenvectors of $\mathbf{A}^H\,\mathbf{A}$. The columns of $\mathbf{U}$, which are the so-called *left singular vectors* of $\mathbf{A}$, are related to the columns of $\mathbf{V}$ by

$$\mathbf{u}^{(i)} = \frac{1}{\sigma_i}\mathbf{A}\,\mathbf{v}^{(i)} \quad , \quad \sigma_i \neq 0 \quad .$$

The related, so-called *Reduced* SVD of $\mathbf{A}$ is

$$\mathbf{A} = \hat{\mathbf{U}}\hat{\boldsymbol{\Sigma}}\;\hat{\mathbf{V}}^H$$

where $\hat{\boldsymbol{\Sigma}}$ is an $r \times r$ diagonal matrix of only the **non-zero** singular values, $\hat{\mathbf{U}}$ is $m \times r$, and $\hat{\mathbf{V}}$ is $n \times r$. (Both $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ have orthonormal columns. However, in the reduced SVD, this only means $\hat{\mathbf{U}}^H\,\hat{\mathbf{U}} = \mathbf{I}$ and $\hat{\mathbf{V}}^H\,\hat{\mathbf{V}} = \mathbf{I}$. Because they are not square, however, we can generally expect that $\hat{\mathbf{U}}\,\hat{\mathbf{U}}^H \neq \mathbf{I}$ and $\hat{\mathbf{V}}\,\hat{\mathbf{V}}^H \neq \mathbf{I}$.)

The SVD plays a major role in both theoretical and practical matrix analysis and computation. For example, the rank of $\mathbf{A}$ is precisely the number of non-zero singular values. Moreover, in terms of the Euclidean norm:

$$\|\,\mathbf{A}\,\|_2 \;=\; \sigma_1 = \|\,\mathbf{A}\mathbf{v}^{(1)}\,\|_2$$

Moreover, we can write the (reduced) SVD in block matrix form as:

$$\mathbf{A} = \left[\;\mathbf{u}^{(1)} \;\vdots\; \mathbf{u}^{(2)} \;\vdots\; \ldots \;\vdots\; \mathbf{u}^{(r)}\;\right] \begin{bmatrix} \sigma_1 & 0 & \ldots & 0 \\ 0 & \sigma_2 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & \sigma_r \end{bmatrix} \begin{bmatrix} {\mathbf{v}^{(1)}}^H \\ {\mathbf{v}^{(2)}}^H \\ \vdots \\ {\mathbf{v}^{(r)}}^H \end{bmatrix}$$

$$= \left[\;\sigma_1\mathbf{u}^{(1)} \;\vdots\; \sigma_2\mathbf{u}^{(2)} \;\vdots\; \ldots \;\vdots\; \sigma_r\mathbf{u}^{(r)}\;\right] \begin{bmatrix} {\mathbf{v}^{(1)}}^H \\ {\mathbf{v}^{(2)}}^H \\ \vdots \\ {\mathbf{v}^{(r)}}^H \end{bmatrix}$$

But this last expression can be viewed as simply the product of a $1 \times r$ block row vector by a $r \times 1$ column vector, i.e. a block inner product. This produces

$$\mathbf{A} = \sigma_1 \mathbf{u}^{(1)} \mathbf{v}^{(1)H} + \sigma_2 \mathbf{u}^{(2)} \mathbf{v}^{(2)H} + \cdots + \sigma_r \mathbf{u}^{(r)} \mathbf{v}^{(r)H} = \sum_{i=1}^{r} \sigma_i \mathbf{u}^{(i)} \mathbf{v}^{(I)H}$$

which is equivalent to expressing $\mathbf{A}$ as a sum of rank one matrices, each of which has smaller Euclidean norm (i.e. σ_i) than all of its predecessors. Viewed slightly differently, this decomposes $\mathbf{A}$ into a sum of rank one matrices of decreasing energy.

This last observation allows us to use the singular values to create approximations to the original matrix. Specifically, if $\nu \leq r$

$$\mathbf{A}^{(\nu)} = \begin{bmatrix} \mathbf{u}^{(1)} & \vdots & \ldots & \vdots & \mathbf{u}^{(\nu)} \end{bmatrix} \begin{bmatrix} \sigma_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \sigma_\nu \end{bmatrix} \begin{bmatrix} \mathbf{v}^{(1)T} \\ \vdots \; \vdots \; \vdots \\ \mathbf{v}^{(\nu)T} \end{bmatrix} = \sum_{i=1}^{\nu} \sigma_i \mathbf{u}^{(i)} \mathbf{v}^{(i)H}$$

represents a so-called rank ν approximation to $\mathbf{A}$. Moreover, we can show that the Euclidean norm of the remaining terms, i.e.

$$\| \sum_{i=\nu+1}^{r} \sigma_i \mathbf{u}^{(i)} \mathbf{v}^{(i)H} \| = \| \mathbf{A} - \mathbf{A}^{(\nu)} \|_2 = \sigma_{k+1}$$

i.e. the next largest singular value. Approximations built on this idea are commonly used in signal and image processing, data compression, etc.

Name: ______________________________

MA 3046 - Matrix Analysis
Laboratory Number 2
The Singular Value Decomposition (SVD)

1. Link to the laboratory web page and download the file:

svdstuff.mat

to your laboratory directory.

2. Start MATLAB and load the data in **svdstuff.mat**. Specifically examine the 6×4 matrix:

$$\mathbf{a} = \begin{bmatrix} \quad & \quad & \quad & \quad \\ \\ \\ \\ \\ \\ \end{bmatrix}$$

3. Using MATLAB's **help** capability, study the syntax of the **eig( )** command until you feel comfortable executing it. Then give the command

$$[\ \mathbf{veig}\ ,\ \mathbf{d}\ ] = \mathbf{eig(\ a'{*}a)}$$

and record both the eigenvalues for $\mathbf{a}' * \mathbf{a}$(where **a** is the matrix created in part 2):

and the associated eigenvectors

$$\mathbf{veig} = \begin{bmatrix} \quad & & \quad \\ & & \\ & & \end{bmatrix}$$

4. Using MATLAB's **sqrt()** command, find the square roots of the eigenvalues found in part 3 above

5. Determine whether or not the eivenvectors (columns of **veig**) found in part 3 above are or are not orthogonal?

Why should this either have or have not been expected?

6. Using MATLAB's **help** capability, study the syntax of the **svd()** command until you feel comfortable executing it. Then give the command

svd(a)

where **a** is the matrix from part 2, and record the results:

How do these values agree with theory when compared with those from part 5 above.

What, if anything, can you infer about the properties of **a** from these values?

7. Next, give the command

$$[\ \mathbf{u}\ ,\ \mathbf{s}\ ,\ \mathbf{v}\ ]\ =\ \mathbf{svd(\ a\ )}$$

where **a** is the matrix from part 2, and record the results:

$$\mathbf{u} = \left[\quad \right]$$

$$\mathbf{s} = \left[\quad \right]$$

$$\mathbf{u} = \left[\quad \right]$$

8. Verify that, for the matrices $\mathbf{u}$, $\mathbf{s}$ and $\mathbf{v}$ found in part 7:

(i) $\mathbf{a} \;=\; \mathbf{u} * \mathbf{s} * \mathbf{v}'$ to some reasonable order.

(ii) $\mathbf{u}'{*}\mathbf{u} \;=\; \mathbf{I}$ to some reasonable order.

(iii) $\mathbf{v}'{*}\mathbf{v} \;=\; \mathbf{I}$ to some reasonable order.

9. Compare the eigenvectors (**veig**) of $\mathbf{a}' * \mathbf{a}$ as found in part 3 with the left singular vectors ($\mathbf{v}$) of $\mathbf{a}$ as determined in part 7 above. Explain any differences.

10. Use MATLAB to find the Euclidean norm of the matrix **a** from part 2. Then compare this value to the singular values of that same matrix as computed in part 6. Why, or why not, should you have expected this result?

11. Give the MATLAB command

norm(a*v(: , 1))

and record the result.

Why, or why not, should you have expected this value?

12. Compute

$$\mathbf{apar} \;=\; \mathbf{u(:,1\!:\!3)*s(1\!:\!3,1\!:\!3)*v(:,1\!:\!3)'}$$

$$\mathbf{ans} = \left[\begin{array}{c} \\ \\ \\ \\ \\ \\ \\ \\ \\ \\ \end{array} \right]$$

What is this calculation actually doing? Compare this result to the original matrix **a** from part 2 and its singular values as computed in part 6. Why, or why not, should the degree to which they compare have been expected?